REMARKS

Prior to this Amendment, Claims 1-54 were pending in the application. All claims have been rejected. In response, certain claims have been amended, Claim 19 has been canceled, and new claims have been added.

Rejections of Claims 1-6, 8-12, 18-30, 32, 34-38, and 44-54

Claims 1-6, 8-12, 18-30, 32, 34-38, and 44-54, of which claims 1, 27, 53 and 54 are independent, have been rejected under 35 U.S.C. §103(a) as being unpatentable over Chaudhuri et al., U.S. Patent No. 5,926,813 (Chaudhuri I) in view of Chaudhuri et al., U.S. Patent No. 6,223,171 (Chaudhuri II). The inclusion of claims 21, 25, 47 and 51 in this group, in section 2, pages 2-7 of the Office Action, appears to be in error, as these claims (21, 25, 47 and 51) are not discussed in section 2, but are specifically rejected in section 6 of the Office Action, pages 9-10. The Applicants respectfully request clarification.

Claims 1 and 27

Claim 1 has been amended for clarity, and to remove references to "baseline statistics," which are now included in new dependent claim 56. Claim 5 has been amended to depend from new claim 56. No new matter has been added.

The Applicants' invention is a mechanism for finding an optimal set of indexes for a database, based on the database schema, workload information, and an index superset formed by a union of the current index set and a *proposed index set*. By starting with this limited superset of indexes, the claimed invention is able to focus the index verification process from early on.

The proposed index set can be proposed, for example, by specialized tools, such as Oracle Expert, that use extensive rule bases to suggest particular indexes. Also, an experienced administrator can feed in a proposed index set. Starting from this limited superset (the union of the current index set and the proposed index set) saves a significant amount of resources and time to reach an optimal solution.

Chaudhuri I, on the other hand, starts with the set of *all possible* indexes, and then attempts to reduce the number of indexes and therefore index configurations for evaluation

(Chaudhuri I, col. 6, lines 62-67). Thus, Chaudhuri I's technique consumes more resources and time to reach an optimal solution than the Applicants' claimed invention.

Chaudhuri I, col. 7, lines 56-63, cited by the Examiner, discusses gathering statistics and storing index entries – but does not teach or suggest forming an index superset from a union of a current index set and a proposed index set, as recited in Applicants' Claim 1.

Applicants therefore respectfully request reconsideration and withdrawal of the rejection of Claim 1 as amended, in favor of allowance. Allowance of dependent claims 2-18 and 20-26, and new dependent claims 55 and 56, follows from the allowance of Claim 1.

Claim 27 has been amended to recite the proposed index set and should now be allowable for the same reasons as Claim 1. Allowance of dependent claims 28-52 and new dependent claim 57 follows from the allowance of Claim 27.

Claims 53 and 54, and new Claims 58-109

Like Claim 27, Claims 53 and 54 have been amended to recite the proposed index set and should now be allowable for the same reasons as Claim 1.

New dependent Claims 58-83 and 84-109, corresponding to Claims 2-18, 20-26, 55 and 56, have been added, and their allowance follows from the allowance of Claims 53 and 54 respectively.

Claims 6 and 32

Claims 6 and 32 have been amended to emphasize that at least one statistic is generated based on an execution plan generated by an optimizer. Claims 7, 8, 11 and 12 have been amended for compatibility with these changes.

An optimizer typically generates a cost, and an "execution plan," which is an optimal series of steps to be executed to fulfill execution of the subject workload. The Applicants' invention, as recited in claims 6 and 32 as amended, evaluates the plan to generate at least one statistic.

For example, an execution plan for a given workload might indicate that table 1 is to be accessed using index a, that table 2 is to be accessed using index b and finally that x rows of information are to be returned. The indicated index utilization, for example, the fact that the

plan uses index a, can contribute to the index usage statistics of the invention as recited in claim 6.

The reason for this is that cost is not sufficient when determining tradeoffs as to whether or not to include an index in the solution. The Applicants have found that consideration of other factors that can be gleaned from the execution plan, such as SQL statement frequency or index volatility, or even a user-specified SQL statement importance, can benefit the index selection process. These additional factors allow the mechanism of the claimed invention to make better choices for indexing solutions when, for example, the optimizer costs are similar.

Furthermore, with a high insert-intensive workload, the Applicants' invention can use the volatility information to guide the solution towards fewer indexes, possibly choosing a solution with a higher cost that minimizes the number of indexes.

Chaudhuri I, at col. 10, lines 22-25 and 53-60, cited by the Examiner, teach that an optimizer returns a cost and an execution plan. The cost is evaluated, but nothing in the passage, alone or in combination with the other references, teaches or suggests that the cost evaluation tool looks at and evaluates the plan itself.

The Applicants therefore respectfully request the reconsideration and withdrawal of the rejections of claims 6 and 32 in favor of allowance, regardless of the final disposition of their respective base claims. Allowance of their respective dependent claims, *i.e.*, claims 7-19 and 33-45, would then follow.

Rejections of Claims 21 and 47

Dependent claims 21 and 47 have been rejected under 35 U.S.C. §103(a) as being unpatentable over Chaudhuri I and Chaudhuri II, in view of Celis et al., U.S. Patent No. 6,021,405). The Applicants respectfully traverse this rejection.

First, the Applicants disagree that it would be obvious to combine Chaudhuri with Celis. Celis is directed to optimizing the execution of a query by transforming logical expressions into execution plans (Celis, Abstract). This is typically done automatically after submission of the query. See, for example, Celis, column 12, lines 40-65. Both Chaudhuri I and Chaudhuri II, on the other hand, are directed to reducing costs and time in selecting index sets (Chaudhuri I, Abstract). Indexes are simulated but not actually created (Chaudhuri I,

column 6, lines 46-50). Thus, as pertains to Chaudhuri, queries are not executed, but rather are used in the analysis of various candidate index sets. The Applicants do not believe that one skilled in the art would be motivated to modify Celis's transformance of logical expressions with Chaudhuri's simulated "what-if" indexes, or vice versa. The Applicants thus respectfully request the reconsideration and withdrawal of the rejections of claim 21 and 47 under 35 U.S.C. §103(a).

Even when combined, however, Chaudhuri I, Chaudhuri II and Celis together do not teach the Applicants' claimed invention.

The Applicants' claimed invention uses a reduced workload of **unique statements** rather than a full workload during the evaluation process (Specfication, page 3, lines 4-5, and Fig. 3A). The reduced workload compresses the set of SQL statements in the workload down to a smaller set that is relevant for index analysis. The Applicants have found that using a full, unreduced workload is not practical for real world applications. Such analysis has been found to be very resource-intensive and time-consuming without the use of a reduced workload.

Celis, col. 5, lines 37-44, on the other hand, discusses tracking rules to eliminate the *generation* of redundant expressions that can occur when rules (i.e., transformation rules which produce equivalent expressions – see Celis, col. 2, lines 25-28) are applied multiple times to an expression. Celis says nothing about *reducing a workload* into **unique statements**, as recited in claims 21 and 47.

That is, Celis does not teach or suggest "the workload is reduced into unique statements," as recited in claims 21 and 47 as filed. Therefore, claims 21 and 47 are allowable regardless of the final dispositions of their respective base claims.

Rejections of Claims 25 and 51

Dependent claims 25 and 51, each of which recites that the proposed index list is provided by an expert system, have been rejected under 35 U.S.C. 103(a) as being unpatentable over Chaudhuri I and Chaudhuri II, in view of Lomet et al. The Applicants respectfully disagree.

"Expert system" is a well-known term in the computer industry. An expert system is

> a computer program that simulates the judgement and behavior of a human or an organization that has expert knowledge and experience in a particular field. Typically, such a system contains a knowledge base containing accumulated experience and a set of rules for applying the knowledge base to each particular situation that is described to the program. Sophisticated expert systems can be enhanced with additions to the knowledge base or to the set of rules.

(http://whatis.techtarget.com/definition/0,,sid9_gci212087,00.html, last updated on July 30, 2001).

Although Lomet, as the Examiner suggests, discusses the use of expert systems to select an index, Lomet actually teaches away from a combination of an expert system with a system such as the Applicants' in which a proposed set of indexes is suggested for further evaluation. For example, in discussing prior work, Lomet states: "The class of tools that adopt an expert system like approach ... *suffer* from being disconnected from the query optimizer." Lomet, page 83, left column. Lomet's tool, the "index tuning wizard," merely iterates through the space of hypothetical indexes. There is no teaching or suggestion that an expert system could be used to provide a proposed index set from which, along with a current index set, statistics can be generated, as recited in claims 25 and 51.

Chaurhuri I, col. 5, line 57 to col. 6, line 13, cited by the Examiner, discusses an index selection tool but does not suggest that this tool makes use of an expert system.

Therefore, claims 25 and 51 are allowable regardless of the final disposition of their respective base claims. The Applicants respectfully request the reconsideration and withdrawal of the rejections of these claims in favor of allowance.

Rejection of Claim 31

Dependent claim 31, which recites "eliminating at least one index on a small table," has been rejected under 35 U.S.C. §103(a) as being unpatentable over Chaudhuri I, Chaudhuri II, Jakobsson and Siegal et al. The Applicants respectfully disagree.

Jakobsson, col. 12, line 66 to col. 13, line 1, cited by the Examiner, teaches that the cost of searching a table without using an index is proportional to the size of the table. The purpose of the discussion appears to be to show how the cost improves when an index is added. However, Jakobsson does not teach or suggest, as the Examiner asserts, *eliminating* an index

on a small table, for example, where "the table is sufficiently small that an index would never be a desirable access method." (Specification as filed, page 20, lines 9-10).

None of the cited references, alone or in combination, teach or suggest "eliminating at least one index on a small table," as recited in claim 31. Therefore claim 31 is allowable. The Applicants therefore respectfully request the reconsideration and withdrawal of the rejection of claim 31, in favor of allowance.

New claims

In addition to the new dependent claims discussed above, new Claims 55, 57 and 110-136 have been added to the application.

Support for new claims 55 and 57 can be found throughout the Specification as filed. (See, *e.g.*, page 13, lines 24-26). No new matter has been added. Allowance of these new claims is respectfully requested regardless of the final disposition of their respective base claims.

New independent Claim 110 parallels the language of Claim 1 as amended, using means-plus-function language and should be allowable for the same reasons as Claim 1. Allowance of new dependent Claims 111-136, corresponding to Claims 2-18, 20-26, 55 and 56 follows from the allowance of Claim 110.

CONCLUSION

In view of the above amendments and remarks, it is believed that all claims are in conditions for allowance, and it is respectfully requested that the application be passed to issue. If the Examiner feels that a telephone conference would expedite prosecution of this case, the Examiner is invited to call the undersigned attorney at (978) 341-0036.

Respectfully submitted,

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

By Gerald M. Bluhm

Gerald M. Bluhm
Registration No.: 44,035
Telephone: (978) 341-0036
Facsimile: (978) 341-0136

Concord, MA 01742-9133
Dated: 2/27/03

MARKED UP VERSION OF AMENDMENTS

Claim Amendments Under 37 C.F.R. § 1.121©(1)(ii)

1. (Amended) A method for evaluating a plurality of candidate index sets for a workload of database statements in a database system, the method comprising:

[generating baseline statistics for each statement in the workload;]

forming an index superset from a union of a current index set and a proposed index set;

deriving a candidate index set from the index superset, the derived candidate index set being included in [one of] the plurality of candidate index sets;

generating statistics based on the derived candidate index set [and the baseline statistics]; and

presenting the generated statistics.

5. (Amended) The method of Claim [1] 56, wherein deriving the baseline statistics comprises disabling current indexes.

6. (Amended) The method of Claim 1, wherein generating statistics for a statement comprises generating at least one statistic based on an execution plan created by an optimizer.

[creating an execution plan which represents a series of steps for executing the statement;

evaluating the execution plan;

generating and recording statistics based on the evaluation of the execution plan].

7. (Amended) The method of Claim 6, wherein [creating an] the execution plan is based on available access paths.

8. (Amended) The method of Claim 6, wherein [creating an] the execution plan is based on statistics for at least one schema object accessed by the statement.

11. (Amended) The method of Claim 6, [wherein evaluating the execution plan comprises] further comprising:

for a table accessed by a statement under evaluation, using the execution plan to identify [identifying] at least one index that would be used to retrieve data from the table upon an execution of the statement.

12. (Amended) The method of Claim 6, wherein [evaluating the execution plan comprises: determining] the optimizer generates a cost of the execution plan.

27. (Amended) A system for evaluating a plurality of candidate index sets for a workload in a database system, the workload derived from a plurality of statements, the system comprising:

a workload evaluator which evaluates each statement within the workload;

an index solution evaluator which, responsive to the workload evaluator, evaluates each index in a candidate index set with respect to the workload, the candidate index solution being one of the plurality of candidate index sets, each candidate index set derived from an index superset formed by the union of a current index set and a proposed index set;

a solution/rollup evaluator which, responsive to the index solution evaluator, evaluates the candidate index solution; and

a solution refiner which, responsive to the solution/rollup evaluator, generates at least one new candidate index solution.

32. (Amended) The system of Claim 27, wherein [the evaluation created by] the workload evaluator evaluates an execution plan created by an optimizer, the execution plan comprising, for each statement of the workload, [comprises] an execution plan [for each statement] which represents a series of steps for executing the statement, [wherein the workload evaluator evaluates the execution plan, and] the workload evaluator

further generating and recording [generates and records] statistics based on the evaluation of the execution plan.

53. (Amended) A computer program product for evaluating a plurality of candidate index sets for a workload of database statements in a database system, the computer program product comprising a computer usable medium having computer readable code thereon, including program code which:

[generates baseline statistics for each statement in the workload;]

forms an index superset from a union of a current index set and a proposed index set;

repeatedly

derives a candidate index set from the index superset, the derived candidate index set [superset] being included in [one of] the plurality of candidate index sets, and

generates statistics based on the derived candidate index set [and the baseline statistics]; and

presents the generated statistics.

54. (Amended) A computer data signal embodied in a carrier wave for evaluating a plurality of candidate index sets for a workload of database statements in a database system, comprising:

[program code for generating baseline statistics for each statement in the workload;]

program code for forming an index superset from a union of a current index set and a proposed index set;

program code for repeatedly

deriving a candidate index set from the index superset, the derived candidate index set [superset] being [one of] included in the plurality of candidate index sets, and

generating statistics based on the derived candidate index set [and the baseline statistics]; and

program code for presenting the generated statistics.